

1. A method for obtaining update information in a disk storage facility which stores data groups as a plurality of data blocks, said method comprising the steps of:
 - A) defining a data group set of at least one data group,
 - B) defining an interval during which the data group is to be monitored,
 - C) recording, during the defined interval, a first update to each data block in the data group set, and
 - D) after the defined interval determining, from said recording, the total number of data blocks in the data group set that were updated at least one time during the defined interval.
2. A method as recited in claim 1 wherein said step of recording includes defining a flag for each data block and altering the state of each flag upon receiving an update to a corresponding data block.
3. A method as recited in claim 2 wherein said step of recording additionally includes:
 - i) establishing a first state for each flag at the beginning of the defined interval, and
 - ii) establishing a second state for an individual flag when the corresponding data block is updated, said recording counting the total

number of flags in the second state at the end of the defined interval.

4. A method as recited in claim 2 wherein said recording step includes storing in an other data set an identification of the time interval, data group and the defined flags.
5. A method as recited in claim 2 wherein the defined interval is divided into subintervals and said recording step includes storing, for each data group in the data group set identifications of the subinterval and data group and the defined flags.
6. A method as recited in claim 5 wherein said step of recording includes:
 - i) establishing a first state for each flag at the beginning of each sample subinterval, and
 - ii) establishing a second state for an individual flag when the corresponding data block is updated, said recording counting the total number of flags in the second state at the end of the sample subinterval.
7. A method as recited in claim 1 wherein updates to the data storage facility are copied over a data communications

5 path having predetermined characteristics to a data facility, said method additionally comprising the step of determining from information obtained during said recording in combination with a first characteristic of the data communications path to obtain a second characteristic of the data communications path.

8. A method as recited in claim 7 wherein the data facility operates as a mirror for the data storage facility for maintaining a synchronized copy of the data therein and wherein the first and second characteristics are, respectively, bandwidth and time for the mirror to achieve synchronism for the recorded update activity for that given bandwidth.

5 9. A method as recited in claim 7 wherein the data facility operates as a mirror for the data storage facility for maintaining a synchronized copy of the data therein and wherein the first and second characteristics are, respectively, the time required for the mirror to achieve synchronism for the recorded update activity and the bandwidth required to achieve that resynchronization time, respectively.

10. A method for determining, from a local site with a local data storage facility, update operations between first and second remote data storage facilities over a communications path wherein each of the data storage facilities stores data in data sets on disk tracks, said method comprising the steps of:

- 5 A) defining a set of disk tracks in the local data storage facility,
- B) defining a time interval,
- 10 C) recording, at the end of the time interval, an identification of each track in the local data storage facility that is updated during the interval, and
- D) upon completion of the time interval determining the number of tracks in the defined set of disk tracks that were updated during the time interval.

15 11. A method as recited in claim 10 wherein said recording includes:

- 5 i) defining a table with a position corresponding to each of the defined disk tracks and each position being set to a first state, and
- ii) responding to an update operation during the defined time interval to a defined disk track by

setting the corresponding position to a second state.

12. A method as recited in claim 11 wherein said recording occurs over a plurality of defined time intervals and said recording sets all the table positions to the first state at the beginning of each interval.
13. A method as recited in claim 12 additionally comprising the step of storing the contents of the table after each interval with a date-time stamp corresponding to the interval.
14. A method as recited in claim 13 wherein said determining step includes:
 - i) defining a report to be generated including a time frame for the report,
 - ii) processing the stored contents of the table with date-time stamps within the time frame, and
 - iii) generating a report with the numbers of track changes recorded in the processed contents.
15. A method as recited in claim 14 wherein the local disk storage facility includes a controller including a set of logical volumes and the defined disk tracks include all

the disk tracks in the controller, said step of defining a report including generating a corresponding reporter command, said processing including combining all the track changes for the controller.

5 16. A method as recited in claim 15 wherein said processing includes combining the date-time stamped contents of each table in the report time frame in a logical OR operation to obtain a final table and counting the total number of positions in the final table set to the second state.

5 17. A method as recited in claim 16 additionally comprising the step of generating the reporter command with one parameter specifying a resynchronization time in which corresponding updates should be completed between the first and second remote data storage facilities, said processing providing the required bandwidth of the communications path to provide that resynchronization time.

5 18. A method as recited in claim 17 wherein the communications path includes at least one path with a characteristic bandwidth, said bandwidth being further determined by dividing the required bandwidth by the number of paths to obtain the characteristic bandwidth.

19. A method as recited in claim 17 wherein the communications path will include at least one path with a characteristic bandwidth, said required bandwidth being further determined by dividing the required bandwidth by the characteristic bandwidth to determine the number of required paths.

5

20. A method as recited in claim 16 additionally comprising the step of generating the reporter command with one parameter specifying an available bandwidth of the communications path, said processing providing a resynchronization time during which all updates to the first remote data storage facility will be transferred over the communications path to the second remote data storage facility.

5

21. A method as recited in claim 20 wherein the communications path includes at least one path having a characteristic bandwidth, the bandwidth provided with the command being the product of the number of paths and the characteristic bandwidth.

22. A method as recited in claim 14 wherein the local data storage facility includes a plurality logical volumes and

the defined disk tracks include all the disk tracks in the controller, said step of generating the reporter command including defining a logical volume report for a set of at least one logical volume, said processing combining all the track changes for the defined logical volumes.

5

23. A method as recited in claim 22 wherein said processing includes combining all the contents of the table for a logical volume in a logical OR operation to obtain a final table and summing the positions in the final table that are set to the second state.

- 5
24. A method as recited in claim 16 wherein the local disk storage facility includes a plurality logical volumes and each logical volume includes at least one data set, said step of generating the reporter command defining a data set report with a set of at least one data set, said processing combining all the track changes for the defined data sets.

25. A method as recited in claim 24 wherein said processing additionally includes defining a mask of the tracks forming the defined data sets, generating a final table that is the logical OR of all the tables in the logical volume containing the defined data sets and thereafter

combining the mask and the final table in a logical AND operation.

26. Apparatus for obtaining update information in a disk storage facility which stores data groups as a plurality of data blocks, said apparatus comprising:

- A) means for defining a data group set of at least one data group,
- B) means for defining an interval during which the data group is to be monitored,
- C) a collector application that records, during the defined interval, a first update to each data block in the data group set, and
- D) a reporter application, activated after the defined interval, that determines, from the recording produced by said collection module, the total number of data blocks in the data group set that were updated at least one time during the defined interval.

10

15

27. Apparatus as recited in claim 26 wherein said collector application includes means for defining a flag for each data block, the state of each flag being altered upon receiving an update to a corresponding data block.

28. Apparatus as recited in claim 27 wherein said collector application includes:
 - i) means for establishing a first state for each flag at the beginning of the defined interval.
29. Apparatus as recited in claim 27 wherein said collector application includes means for storing in an other data set an identification of the time interval, data group and all the defined flags.
30. Apparatus as recited in claim 27 wherein the defined interval is divided into subintervals and said collector application includes means for storing, for each data group in the data group set, identifications of the subinterval and data group and the defined flags.
5
31. Apparatus as recited in claim 30 wherein said collector application includes means for establishing a first state for each flag at the beginning of each sample subinterval.
32. Apparatus as recited in claim 26 wherein updates to the data storage facility are copied over a data communications path having predetermined characteristics to a data facility, said apparatus additionally comprising a reporter application that processes the data set
5

produced by said collector application in combination with a first characteristic of the communications path to obtain a second characteristic of the communications path.

5

33. Apparatus as recited in claim 26 wherein the data facility operates as a mirror for the data storage facility for maintaining a synchronized copy of the data therein, said apparatus additionally comprising a reporter application that process the data set produced by said collector application, first and second characteristics are, respectively, bandwidth and time for the mirror to achieve synchronism for the recorded update activity.
34. Apparatus as recited in claim 26 wherein the data facility operates as a mirror for the data storage facility for maintaining a synchronized copy of the data therein and wherein the first and second characteristics are, respectively, the time required for the mirror to achieve synchronism for the recorded update activity and the bandwidth required to achieve that resynchronization time, respectively.
35. Apparatus for determining, from a local site with a local data storage facility, update operations between first and second remote data storage facilities over a

communications path wherein each of the data storage facilities stores data in data sets on disk tracks, said method comprising the steps of:

- A) means for defining a set of disk tracks in the local data storage facility,
- B) means for defining a time interval,
- C) a collector application for recording, at the end of the time interval, an identification of each track in the local data storage facility that is updated during the interval, and
- D) a reported application for determining the number of tracks in the defined set of disk tracks that were updated during the time interval.

36. Apparatus as recited in claim 35 wherein said collector application includes means for defining a table with a position corresponding to each of the defined disk tracks and each position being set to a first state, and said apparatus includes means for responding to an update operation during the defined time interval to a defined disk track by setting the corresponding position to a second state.

37. Apparatus as recited in claim 36 wherein said collector application operates over a plurality of defined time

5

intervals and includes means for setting all the table positions to the first state at the beginning of each interval.

38. Apparatus as recited in claim 37 additionally comprising means for storing the contents of the table after each interval with a date-time stamp corresponding to the interval.
39. Apparatus as recited in claim 38 wherein said reported application operates in response to a command defining a report to be generated including a time frame for the report and includes:
 - 5 i) means for processing the stored contents of the table with date-time stamps within the time frame, and
 - ii) means for generating a report with the numbers of track changes recorded in the processed contents.
- 10 40. Apparatus as recited in claim 39 wherein the local disk storage facility includes a controller including a set of logical volumes and the defined disk tracks include all the disk tracks in the controller, said reported application being responsive to a reporter command

5

specifying a controller report and including means for combining all the track changes for the controller.

41. Apparatus as recited in claim 40 wherein said combining means includes means for combining the date-time stamped contents of each table in the report time frame in a logical OR operation to obtain a final table and means for counting the total number of positions in the final table set to the second state.

5

42. Apparatus as recited in claim 41 wherein the reporter command has one parameter specifying a resynchronization time in which corresponding updates should be completed between the first and second remote data storage facilities, said report generating means providing the required bandwidth of the communications path to provide that resynchronization time.

5

43. Apparatus as recited in claim 42 wherein the communications path includes at least one path with a characteristic bandwidth, said report generating means including means for dividing the required bandwidth by the number of paths to obtain the characteristic bandwidth.

44. Apparatus as recited in claim 42 wherein the communications path will include at least one path with a characteristic bandwidth, said report generating means including means for dividing the required bandwidth by the characteristic bandwidth to determine the number of required paths.

5

45. Apparatus as recited in claim 41 wherein the reporter command specifies an available bandwidth of the communications path, said report generating means providing a resynchronization time during which all updates to the first remote data storage facility will be transferred over the communications path to the second remote data storage facility.

5

46. Apparatus as recited in claim 45 wherein the communications path includes at least one path having a characteristic bandwidth, said report generating means including means for generating the product of the number of paths and the characteristic bandwidth.

47. Apparatus as recited in claim 39 wherein the local data storage facility includes a plurality logical volumes and the defined disk tracks include all the disk tracks in the controller, said reporter application responding to a

5 logical volume report command identifying a set of at least one logical volume and said combining means combining all the track changes for the defined logical volumes.

48. Apparatus as recited in claim 47 wherein said combining means includes means for combining all the contents of the table for a logical volume in a logical OR operation to obtain a final table and means for summing the positions in the final table that are set to the second state.

5

49. Apparatus as recited in claim 41 wherein the local disk storage facility includes a plurality logical volumes and each logical volume includes at least one data set, said reported application responding to a command defining a data set report with an identification of a set of at least one data set, said combining means combining all the track changes for the defined data sets.

50. Apparatus as recited in claim 49 wherein said reported application additionally includes means for defining a mask of the tracks forming the set of defined data sets, means for generating a final table that is the logical OR of all the tables in the logical volume containing the

E30-033CON (99-035CON)

defined data sets and means for combining the mask and the final table in a logical AND operation.